Uniform Error Bounds for Gaussian Process Regression with Application to Safe Control


Armin Lederer
armin.lederer@tum.de


Jonas Umlauft
jonas.umlauft@tum.de


Sandra Hirche
hirche@tum.de

Motivating Example

Nonparameteric regresssion offers great promises in robotic applications

Learned policies are unsafe in real world applications [1]

- Constrained environments to avoid damages of hardware
- No human-robot interaction due to risk of injuries

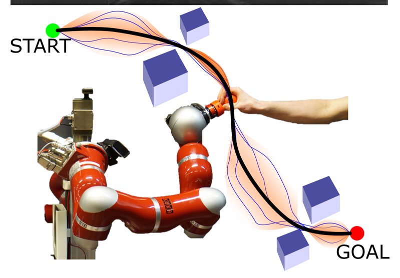

Quantification of uncertainty in data-driven models essential for safety-critical applications
$\Rightarrow$ Robust control for rigorous safety certificates

How can the learning error be bounded based on the model uncertainty?
How are formal safety guarantees provided for policies based on uncertain models?

Gaussian Process Regression

- Bayesian nonparametric modeling as "distribution over functions"

$$f_{\mathcal{GP}}(\boldsymbol{x}) \sim \mathcal{GP}(0, k(\boldsymbol{x}, \boldsymbol{x}'))$$

- Based on training data $\mathbb{D} = \left\{\boldsymbol{x}^{(i)}, y^{(i)} = f\left(\boldsymbol{x}^{(i)}\right) + \epsilon^{(i)}\right\}_{i=1}^{N}$ with Gaussian noise $\epsilon^{(i)} \sim \mathcal{N}(0, \sigma_n^2)$, it provides mean and variance

$$\nu_N(\boldsymbol{x}) := \mathbb{E}\left[f_{\mathcal{GP}}(\boldsymbol{x})|\boldsymbol{x}, \mathcal{D}\right] = \boldsymbol{k}^\intercal(\boldsymbol{K} + \sigma_n^2 \boldsymbol{I}_N)^{-1}\boldsymbol{y})$$
$$\sigma_N^2(\boldsymbol{x}) := \mathbb{V}\left[f_{\mathcal{GP}}(\boldsymbol{x})|\boldsymbol{x}, \mathcal{D}\right] = k(\boldsymbol{x}, \boldsymbol{x}) - \boldsymbol{k}^\intercal(\boldsymbol{K} + \sigma_n^2 \boldsymbol{I}_N)^{-1}\boldsymbol{k}$$

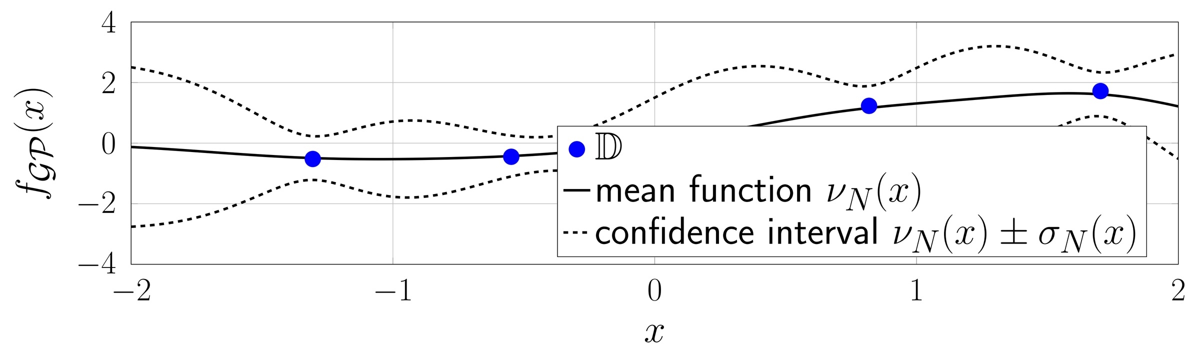

Problem: Difficult quantification of uniform error bounds using RKHS theory [2, 3]

Probabilistic Uniform Error Bound

- Assumption: function $f(\boldsymbol{x})$ is a sample from a GP with Lipschitz constant L_f
- Lipschitz continuous posterior mean $\nu_N(\cdot)$ and standard deviation $\sigma_N(\cdot)$ with

$$\|\nu_N(\boldsymbol{x}) - \nu_N(\boldsymbol{x}')\| \leq L_\nu \|\boldsymbol{x} - \boldsymbol{x}'\| \qquad \|\sigma_N(\boldsymbol{x}) - \sigma_N(\boldsymbol{x}')\| \leq \omega_\sigma(\|\boldsymbol{x} - \boldsymbol{x}'\|)$$

Theorem

The learning error is probabilistically bounded by

$$P\left(|f(\boldsymbol{x}) - \nu_N(\boldsymbol{x})| \leq \sqrt{\beta(\tau)}\sigma_N(\boldsymbol{x}) + (L_\nu + L_f)\tau + \sqrt{\beta(\tau)}\omega_\sigma(\tau),\ \forall \boldsymbol{x} \in \mathbb{X}\right) \geq 1 - \delta$$

with $\beta(\tau) = 2\log((1 + \frac{r}{\tau})^d \delta^{-1})$ on the compact set $\mathbb{X} \subset \mathbb{R}^d$ with maximal extension r for every $\tau \in \mathbb{R}_+$, $\delta \in (0, 1)$.

Probabilistic Lipschitz Constants

- Kernel with continuous partial derivatives up to the fourth order
- Partial derivative kernels

$$k^{\partial i}(\boldsymbol{x}, \boldsymbol{x}') = \frac{\partial^2}{\partial x_i \partial x_i'} k(\boldsymbol{x}, \boldsymbol{x}') \quad \forall i = 1, \ldots, d$$

- Lipschitz constants $L_k^{\partial i}$ of partial derivative kernel

Theorem

The constant

$$L_f = \left\| \begin{bmatrix} \sqrt{2\log\left(\frac{2d}{\delta_L}\right)} \max\limits_{\boldsymbol{x}\in\mathbb{X}} \sqrt{k^{\partial 1}(\boldsymbol{x}, \boldsymbol{x})} + 12\sqrt{6d} \max\left\{ \max\limits_{\boldsymbol{x}\in\mathbb{X}} \sqrt{k^{\partial 1}(\boldsymbol{x}, \boldsymbol{x})}, \sqrt{r L_k^{\partial 1}} \right\} \\ \vdots \\ \sqrt{2\log\left(\frac{2d}{\delta_L}\right)} \max\limits_{\boldsymbol{x}\in\mathbb{X}} \sqrt{k^{\partial d}(\boldsymbol{x}, \boldsymbol{x})} + 12\sqrt{6d} \max\left\{ \max\limits_{\boldsymbol{x}\in\mathbb{X}} \sqrt{k^{\partial d}(\boldsymbol{x}, \boldsymbol{x})}, \sqrt{r L_k^{\partial d}} \right\} \end{bmatrix} \right\|$$

is a Lipschitz constant of $f(\cdot)$ on $\mathbb{X}$ with probability of at least $1 - \delta_L$.

Safe Control of Unknown Dynamical Systems

- Nonlinear control affine dynamical system

$$\dot{x}_1 = x_2, \qquad \dot{x}_2 = f(\boldsymbol{x}) + u,$$

- Goal: track reference $x_d(t)$ with x_1 such that error $\boldsymbol{e} = \boldsymbol{x} - [x_d\ \dot{x}_d]^T$ vanishes
- Define filtered state $r = \lambda e_1 + e_2$, $\lambda > 0$
- Use feedback linearizing policy

$$u = \pi(\boldsymbol{x}) = -\hat{f}(\boldsymbol{x}) + \nu$$

with control gain $k_c > 0$ in linear controller

$$\nu = \ddot{x}_d - k_c r - \lambda e_2$$

Theorem

The feedback linearizing controller with $\hat{f}(\cdot) = \nu_N(\cdot)$ guarantees with probability $1 - \delta$ that the tracking error $\boldsymbol{e}$ converges to

$$\mathbb{B} = \left\{ \boldsymbol{x} \in \mathbb{X} \,\middle|\, \|\boldsymbol{e}\| \leq \frac{\sqrt{\beta(\tau)}\sigma_N(\boldsymbol{x}) + (L_\nu + L_f)\tau + \sqrt{\beta(\tau)}\omega_\sigma(\tau)}{k_c\sqrt{\lambda^2 + 1}} \right\}.$$

Numerical Evaluation on a Robotic Manipulator

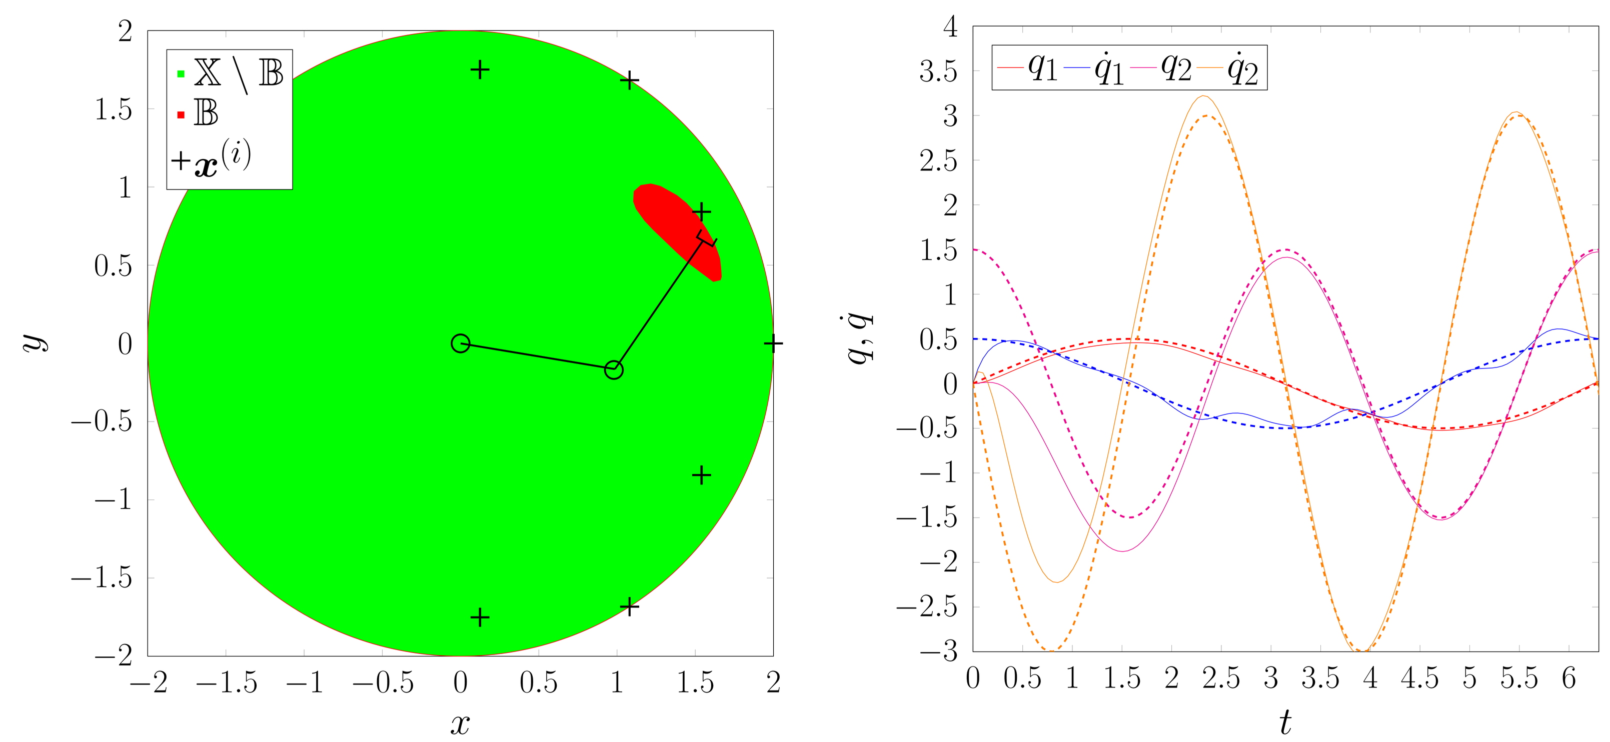

References

[1] E. T. Campolettano, M. L. Bland, R. A. Gellner, D. W. Sproule, B. Rowson, A. M. Tyson, S. M. Duma, and S. Rowson, "Ranges of injury risk associated with impact from unmanned aircraft systems," *Annals of Biomedical Engineering*, vol. 45, 2017.

[2] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger, "Information-theoretic regret bounds for Gaussian process optimization in the bandit setting," *IEEE Transactions on Information Theory*, vol. 58, no. 5, pp. 3250–3265, 2012.

[3] S. R. Chowdhury and A. Gopalan, "On Kernelized Multi-armed Bandits," in *Proceedings of the International Conference on Machine Learning*, 2017, pp. 844–853.

Technische Universität München
Chair of Information-oriented Control